

TITLE OF THE INVENTION

Method for Traffic Engineering
in a Multi-homed Virtual Private Local Area Network Service

CROSS-REFERENCES TO RELATED APPLICATIONS

[0001] Not Applicable.

5

STATEMENT REGARDING FEDERALLY SPONSORED RESEARCH OR
DEVELOPMENT

[0002] Not Applicable.

10

BACKGROUND OF THE INVENTION

[0003] The present embodiments relate to computer networks and are more particularly directed to a method for traffic engineering in a multi-homed Virtual Private Local Area Network Service of a Metro Ethernet Network.

15 [0004] Ethernet networks have found favor in many applications in the networking industry for various reasons. For example, Ethernet is a widely used and cost effective medium, with numerous interfaces and capable of communications and various speeds up to the Gbps range. Ethernet networks may be used to form a Metro Ethernet Network ("MEN"), which is generally a publicly accessible network that provides a Metro domain,
20 typically under the control of a single administrator, such as an Internet Service Provider

("ISP"). A MEN is typically used to connect between an access network and a core network. The access network typically includes private or end users making connectivity to the network. The core network is used to connect to other Metro Ethernet Networks, and the core network provides primarily a packet switching function.

5 [0005] A MEN typically consists of a number of Provider Edge ("PE") nodes that are identified and configured for communicating with one another prior to the communication of packet traffic. The PE nodes are connected in a point-to-point manner, that is, each PE node is connected to another PE node in an emulated and bi-directional virtual circuit manner, where each such connection is achieved by a Label Switched Path
10 ("LSP"). An LSP is sometimes informally referred to as a link. Thus, each PE node may communicate to, and receive packets from, an adjacent PE node. Further, along each LSP, between adjacent PE nodes, are often a number of Provider ("P") nodes. The P nodes maintain no state information and serve primarily a routing function and, thus, are understood not to disturb the point-to-point connection between the PE nodes of the
15 MEN, which are more intelligent devices. A different number of P nodes may be connected in one communication direction between two adjacent PE nodes as compared to the reverse communication direction between those same two adjacent PE nodes.

[0006] PE nodes in the MEN are also coupled to one or more Customer Edge ("CE") nodes, where those CE nodes thereby represent the interface between the MEN and an
20 adjacent access network. Often in the art, the coupling of a PE node to a CE node is through an intermediate node between the PE node and the CE node, where such an intermediate node is referred to as a layer 2 Provider Edge ("L2PE") node. The connectivity between an L2PE node and additional PE nodes in the MEN is typically referred to as homing. More particularly, if the L2PE node is connected to a single PE
25 node, then the connection is referred to as a single home; in contrast, if the L2PE node is connected to more than one PE node, then the connection is referred to as multi-homed. Further, while L2PE nodes may be multi-homed (i.e., connected to more than one other PE node), the CE nodes of any adjacent access network may not, that is, each CE node may connect to only a single L2PE node.

[0007] With the development of the MEN architecture, there have further evolved additional topologies associated with such a network. One example, that pertains to the preferred embodiments that are described later, is the virtual private local area network service ("VPLS"). A VPLS creates an emulated local area network ("LAN") segment for a given set of nodes in a MEN. The VPLS delivers an ISO layer 2 broadcast domain that is fully capable of learning and forwarding on Ethernet MAC addresses that is closed to a given set of nodes. Thus, within the VPLS, packets may be broadcast to all nodes on the VPLS. VPLS also may be included within the above-described framework that includes PE and L2PE nodes, subject to various constraints. First, more than one VPLS may be included in a single MEN and, thus, certain PE nodes of that MEN may be a part of more than one VPLS. Second, with a multiple VPLS MEN, an L2PE node may support more than one VPLS, where each such VPLS has its own respective homing, that is, for each VPLS, that L2PE has a connection to one (and only one) PE node in the MEN.

[0008] Given the various nodes, attributes, and connectivity described above and known in the art, complexities arise in traffic engineering with such parameters, that is, in establishing network communications, appropriate numbers of VPLSs, connectivity, and efficient use of bandwidth. These complexities arise both in establishing these parameters in a new network for the first time as well as modifying that network if one or more factors change over time, such as when a new VPLS is added. These complexities are further complicated by the desire to include 1+1 protection in a network, whereby a first set of parameters are provided, sometimes referred to as a primary network, but are supplemented by a second set of parameters, sometimes referred to as a secondary or backup network, to operate should the first network become inoperable.

[0009] In view of the above, the preferred embodiments provide a device (e.g., network node) with sufficient processing functionality and programmed to provide traffic engineering of a VPLS network with multi-homing, with unicast and multicast traffic, and with 1+1 protection, as further detailed below.

BRIEF SUMMARY OF THE INVENTION

[0010] In the preferred embodiment, there is a processing device programmed to determine homing paths for a plurality of virtual private local area network services in an Ethernet network comprising a plurality of PE nodes. The processing device is
5 programmed to perform the steps of computing a plurality of sets of different homing configurations, computing a cost function for each set of different homing configurations in the plurality of sets of different homing configurations, and selecting a set of homing configurations from the plurality of sets of different homing configurations in response to a respective computed cost function. Each homing configuration in each set of different
10 homing configurations is computed by a respective iteration of steps, where each iteration corresponds to a respective virtual private local area network service in the plurality of virtual private local area network services and for a respective selected layer two provider edge node in the Ethernet network. Each iteration comprises the steps of selecting an ingress PE node and an egress PE node, determining bandwidth into the ingress PE node,
15 determining bandwidth out of the egress PE node, and specifying a first path for communication from the ingress PE node to the egress PE node and a second path for communication from the egress PE node to the ingress PE node, wherein each path of the first and second paths comprises at least one P node.

[0011] Other aspects are also described and claimed.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWING

[0012] Figure 1 illustrates a network system according to the preferred embodiment and having connectivity which may be configured by a central manager in the network system also according to the preferred embodiment.

- 5 [0013] Figure 2 illustrates a flow chart of the preferred embodiment methodology of the central manager to determine the VPLS connectivity of the network.

DETAILED DESCRIPTION OF THE INVENTION

[0014] By way of illustration of one preferred inventive implementation, Figure 1 depicts a network system designated generally at 10. Network system 10, in the preferred embodiments, includes a Metro Ethernet Network ("MEN") that includes a virtual private
5 local area network service ("VPLS"). System 10 as shown in Figure 1 illustrates such a network with various aspects that are generally known in the art, where the illustration is included to introduce various concepts and conventions and also because it is established and modified according to the preferred embodiments as detailed later.

[0015] By way of example, the MEN of system 10 includes five Provider Edge ("PE")
10 nodes PE₁, PE₂, PE₃, PE₄, and PE₅, where the choice of five is only as an illustration and one skilled in the art should appreciate that any number of such nodes may be included. Indeed, note that often a MEN will include considerably more than five PE nodes. Each PE node may be constructed as a processing device by one skilled in the art using various hardware, software, and programming so as to perform the functionality described in this
15 document as well as that known in the art. Further, as a MEN system, while not shown but as also mentioned earlier in the Background Of The Invention section of this document, it should be understood that between adjacent PE nodes there may be located a number of Provider ("P") nodes. In the MEN of system 10, preferably the network is fully meshed, that is, each PE node PE_x is connected to every other PE node in the system,
20 where each connection is by way of a respective Label Switched Path ("LSP") and is illustrated in Figure 1 with an arrow. Further, for sake of reference in this document, the bi-directional connection between two nodes is by way of two opposite direction LSPs, that is, by an LSP pair ("LSPP") that includes one LSP for communications in one direction from a first PE node to a second node and another LSP for communications in the opposite
25 direction, that is, from the second PE node to the first PE node. As an example, PE node PE₁ is connected, via four respective LSPPs, to each of PE nodes PE₂, PE₃, PE₄, and PE₅. Also, for sake of convention, the label used for each LSPP in Figure 1 identifies the two PE nodes between which the LSPP is connected. For example, between PE nodes PE₁ and PE₂ is an LSPP_{1A2}; as another example, between PE nodes PE₃ and PE₄ is an LSPP_{3A4}. To

simplify Figure 1, only a few of these LSPPs are so labeled. Given the connectivity of system 10, any PE node may communicate as a source directly along an LSP to any other PE node as a destination, where that destination PE node may respond along another LSP (albeit through a different set of P nodes) in the reverse direction back to the source PE node. A single communication between two PE nodes in one direction and in this manner is referred to in the art as a unicast communication and, thus, the various LSPPs of system 10 in Figure 1 define the point-to-point interfaces along which unicast traffic is allowed to pass. When a single PE node endeavors to communicate a packet to more than one destination PE node, then such a communication by way of contrast is referred to in the art as a multicast communication.

[0016] Continuing with the MEN of system 10 in Figure 1, it also includes a number of layer 2 PE nodes ("L2PE nodes") L2PE₁, L2PE₂, L2PE₃, L2PE₄, and L2PE₅, where the choice of five is only as an illustration and one skilled in the art should appreciate that any number of such nodes may be included. Each L2PE is also connected to at least one PE node in the MEN. More specifically, each L2PE supports one or more logical entities corresponding to one or more VPLSs, and each such VPLS logical entity has a single LSPP connection to only one PE node in the MEN; to simplify the illustration, each such single LSPP connection is shown with only a single dashed line, although it should be understood to include two LSPs in the same manner as the other PE node-to-PE node connections in Figure 1. By way of example and looking to L2PE node L2PE₁, it has a VPLS logical entity VPLS₁ connected to PE node PE₂, a VPLS logical entity VPLS₂ connected to PE node PE₁, and a VPLS logical entity VPLS₃ connected to a PE node PE₅. Each VPLS logical entity is part of an overall VPLS; for example, VPLS logical entity VPLS₁ in L2PE₁ is part of VPLS₁, which is also supported in L2PE₃ and L2PE₄. Note also that L2PE node L2PE₁ is multi-homed, that is, it is connected to more than one PE node; in contrast, L2PE node L2PE₅ is single homed in that it supports only a single VPLS logical entity (i.e., VPLS₃), and it is connected to a single PE node PE₁. Notwithstanding the single homing of L2PE node L2PE₅, network system 10 is considered multi-homed because it includes one or more other multi-homed L2PE nodes. Additionally, by definition, an L2PE node provides an interface between the MEN and an access network

that includes customer nodes, where at this interface each L2PE node is connected to one or more customer edge ("CE") nodes; indeed, typically a single L2PE node is connected to one CE node for each logical VPLS entity in that L2PE node, that is, each CE node is associated with a respective logical VPLS entity in the L2PE node. To illustrate these aspects, and by way of example with respect to L2PE node L2PE₁, it is connected to three CE nodes CE_{1,1}, CE_{1,2}, and CE_{1,3}, corresponding to its three logical VPLS entities VPLS₁, VPLS₂, and VPLS₃. Similar connectivity is shown for all other L2PE nodes L2PE_x in system 10 being connected to multiple respective CE nodes, where by way of example for each such L2PE nodes having *V* logical VPLS entities, then the L2PE node is connected to *V* CE nodes.

[0017] Given the connectivity in Figure 1, one skilled in the art should appreciate that the example depicts three different VPLSs, where the homing between each L2PE and its associated PE nodes are summarized in the following Table 1:

L2PE	Homed PE node(s)
L2PE ₁	PE ₁ , PE ₂ , PE ₅
L2PE ₂	PE ₂
L2PE ₃	PE ₃ , PE ₄
L2PE ₄	PE ₄ , PE ₅
L2PE ₅	PE ₁

Table 1

From Table 1 and Figure 1, note also that different logical VPLS entities in a same L2PE node may connect to a same PE node, as is shown by way of example with respect to L2PE node L2PE₄, which supports both a logical VPLS entity VPLS₂ and a logical VPLS entity VPLS₃, both of which area connected to a same PE node PE₄. Moreover, the collection of PE nodes connected to each VPLS may be summarized, as shown in the following Table 2:

VPLS	Homed PE node(s)
VPLS ₁	PE ₂ , PE ₃ , PE ₅
VPLS ₂	PE ₁ , PE ₄
VPLS ₃	PE ₅ , PE ₂ , PE ₄ , PE ₁

Table 2

Thus, Table 2 identifies the PE nodes that belong to each VPLS; for example, only PE nodes PE₂, PE₃, and PE₅ belong to VPLS₁, and so forth for the other two VPLSs with their respective PE nodes. Thus, given the illustration of Figure 1 and Tables 1 and 2, one skilled in the art should appreciate the three different VPLSs and the PE nodes connected among each such VPLS.

[0018] Concluding the description of Figure 1, note also that it illustrates only a first set of connections for a first, or primary, network. In other words, for each logical VPLS entity supported by a respective L2PE node, that logical VPLS entity is shown as connected to only a single PE node; however, it should be understood that Figure 1 may be redrawn with different connectivity to illustrate a second, or backup, network whereby each logical VPLS entity, supported by a respective L2PE, is connected to a different single PE node as compared to that shown in Figure 1. The combination of the primary and backup networks will be understood by one skilled in the art to thereby provide so-called 1+1 protection so that if a connection failure occurs on one network (e.g., the first or "primary" network), then traffic may be established as an alternative along the second, or backup, network. Lastly, note that PE node PE₁ is shown in Figure 1 to include a block depicting a central manager CM. In the preferred embodiment, central manager CM is intended to represent an overall processing function that develops the traffic engineering in the multi-homed network system 10. This function may be included in one of the PE nodes such as illustrated in Figure 1, or alternatively a central manager CM may be a separate node or computing device that is not one of the PE nodes in the network. In either case, central manager CM represents a processing device with sufficient knowledge of the state of the network to which it is coupled so as to provide the remaining functionality described herein. The actual hardware, software, and programming to

achieve such a device are readily ascertainable by one skilled in the art given the functional description in this document.

[0019] Figure 2 illustrates a flowchart of a method 20 of steps performed by central manager CM of Figure 1 in its effort to conduct traffic engineering as introduced above, that is, central manager CM is programmed to perform these steps with the results described later. By way of additional background to method 20, note generally that method 20 operates in an iterative fashion to optimize a cost function that, when optimized, provides primary and backup homes in network system 10, in addition to the LSPs. In general, the preferred embodiment choice of homes for the various VPLSs, and the design of the primary and backup paths, is such that traffic is optimized across the various PE and P nodes and links in the network, and 1+1 protection is provided in the event of node or link failure. In addition, unicast traffic and multicast traffic preferably is considered. Also, the preferred embodiment may operate with respect to either a decoupled VPLS or a hierarchical VPLS. In a decoupled VPLS, multi-homing is supported, with one home for each VPLS within an L2PE node L2PE_x, as shown in Figure 1. In a hierarchical VPLS, only dual homing is supported, with one home being the primary home, and the other home being the backup home. In any event, for sake of simplification, the following discussion assumes a decoupled VPLS such as shown in Figure 1, where one skilled in the art will appreciate the comparable aspects as applying to hierarchical VPLS. Additionally, for more details with respect to decoupled VPLS, the reader is referred to K.Kompella et.al, Decoupled Virtual Private LAN Services, IETF draft, "draftkompella-ppvnp-dtls-02.txt", which is hereby incorporated by reference. Similarly, for more details with respect to hierarchical VPLS, the reader is referred to M.Lasserre, V.Kompella et.al, Virtual Private LAN Services over MPLS, IETF draft, "draft-lasserre-vkompella-ppvnp-vpls-03.txt", which is hereby incorporated by reference.

[0020] Turning now to method 20 in more detail, it commences with a step 30. In step 30, central manager CM learns the MEN physical topology of network system 10 as well as the available bandwidth of system 10. This step 30, therefore, is intended to represent that various topology and bandwidth information is made available to central manager

CM. The manner of making such information known to central manager CM may be automated in certain respects, such as by having the various nodes within system 10 provide signaling to central manager CM to provide the topology information. Alternatively, some or all of this information may be manually provided to central manager CM. Further, where method 20 is performed to update the traffic engineering on a previously-configured network, then some of the topology information may be communicated electronically from one or more nodes in network system 10. In any event, in the preferred embodiment, the topology information includes the set of all P nodes, all PE nodes, and all L2PE nodes in the MEN; thus, in the example of Figure 1, the topology information includes PE nodes PE_1 through PE_5 , the P nodes (not shown but described earlier), and the L2PE nodes $L2PE_1$ through $L2PE_5$. In addition, the topology information preferably includes all physical links, where a link, also referred to as an edge, is edge $e=(a,b)$, with a or b belonging to all nodes and thus each edge being a connection between two such nodes. In addition, central manager CM is informed of the network bandwidth given the physical topology. More particularly, the bandwidth includes the bandwidth capability between each adjacent P node P_i and P_j inside the MEN. In addition, the bandwidth includes the bandwidth capability between each L2PE node $L2PE_i$ node and each PE node PE_j . Lastly, the bandwidth includes the expected traffic between various CEs for various VPLSs, that is, for all CEs belonging to a VPLS, the traffic going from CE node CE_i to CE_j is given. Note that in the preferred embodiment this expected traffic may be assembled in the form of a two-dimensional matrix that includes unicast and multicast traffic bandwidth expectations; for example, such a matrix for unicast and multicast traffic may have the information as shown in the following Table 3.

	egress CE_1	egress CE_2	*****	egress CE_n
ingress CE_1	0	1.4 Gb	2 Gb	1.7 Gb
ingress CE_2	3.1 Gb	0	2.3 Gb	2.5 Gb
*	2.9 Gb	1.7 Gb	0	2.4 Gb
*				
*				
ingress CE_n	3 Gb	2.3 Gb	2.8 Gb	0

Table 3

In Table 3, the terms ingress and egress are intended to identify the direction of traffic for a given CE node as it communicates traffic with respect to the MEN. Thus, an ingress CE node CE_n is one that is receiving traffic which will then enter into the MEN, and an egress CE node CE_n is one that is receiving traffic that is exiting from the MEN. Also in step 30, and in relation to the bandwidth information as relating to a CE node and its associated VPLS, central manager CM is informed of the desired VPLSs for each L2PE node; as demonstrated later, however, this desired result may not necessarily be achieved in an optimal connectivity configuration, but this information is available and indeed provides in part a baseline from which the preferred embodiment determines the connectivity within network system 10. In any event, once the physical topology, bandwidth, and VPLS information is available to central manager CM in step 30, then method 20 continues from step 30 to step 32.

[0021] Starting with and following step 32, central manager CM performs various operations with respect to a given VPLS node of network system 10, where in an exemplary embodiment the operations are first performed for a first VPLS at a first L2PE node, followed by operations for the first VPLS at a second L2PE node, and so forth, until all L2PE nodes are comprehended. Thereafter, the steps are repeated for a second VPLS at the first L2PE node, followed by operations for the second VPLS at the second L2PE node, and so forth, until eventually therefore various connectivity configurations are investigated for all desired VPLSs across L2PE nodes of the MEN, as will be further apparent from the remaining discussion.

[0022] Turning now to step 32 in more detail, central manager CM performs what may be perceived as three general operations, with respect to a given VPLS. Considering then a given VPLS (e.g., $VPLS_1$ in Figure 1), and with respect to a first L2PE node (e.g., L2PE node $L2PE_1$ in Figure 1), the first of the three step 32 operations is the selection of a primary ingress node, hereafter indicated as PE node PE_{pi} and a primary egress PE node, hereafter indicated as PE node PE_{pe} , where in both cases the reference to these PE nodes as "primary" is intended to demonstrate that they are to be considered for the primary portion of the 1+1 configuration, as opposed to the backup (or secondary) portion of that

same configuration. Thus, in the example of Figure 1, this first operation may select PE node PE_2 as PE node PE_{pi} and PE node PE_5 as PE node PE_{pe} . In a second operation, central manager CM determines the bandwidth to be provided into the primary ingress node PE_{pi} ; in the preferred embodiment, this determination may be made with reference to the two-dimensional bandwidth table identified in step 30, as shown by way of example in Table 3, since that table provides CE node-to-CE node bandwidth, and that connectivity is through the intermediate PE nodes in system 10. In a third operation, central manager CM determines the bandwidth to be provided out of the primary egress node PE_{pe} ; in the preferred embodiment, this determination also may be made with reference to the two-dimensional bandwidth table identified in step 30, as shown by way of example in Table 3. Next, method 20 continues from step 32 to step 34.

[0023] In step 34, central manager CM operates in a comparable manner as it did in step 32, with a difference being that the PE nodes at issue in step 34 are for the backup (or secondary) connectivity in the 1+1 protection configuration. Thus, in step 34, again there are what may be perceived as three operations. First, central manager CM selects a secondary ingress node, hereafter indicated as PE node PE_{si} and a secondary egress PE node, hereafter indicated as PE node PE_{se} , where in both cases the reference to these PE nodes as "secondary" is intended to demonstrate that they are to be considered for the secondary portion of the 1+1 configuration. In a second operation, central manager CM determines the bandwidth to be provided into the secondary ingress node PE_{si} , and in a third operation, central manager CM determines the bandwidth to be provided out of the secondary egress node PE_{se} . Recall that step 32, above, determined bandwidth into and out from PE nodes PE_{pi} and PE_{pe} , respectively; in step 34, these same bandwidth values are used for PE nodes PE_{si} and PE_{se} , respectively, in order to provide the same bandwidth for the primary and secondary routes. Next, method 20 continues from step 34 to step 36.

[0024] In step 36, central manager CM establishes LSPPs between PE nodes PE_{pi} and PE_{pe} , that is, for the primary route. In the preferred embodiment, the selection of the LSPPs is by selecting the P nodes to form those LSPPs, and as seen later these selections are also limited by various constraints. Further, the criteria for the step 36 selection are

preferably based on either distribution of the load among all P nodes in the LSPP or, alternatively, the shortest number of P nodes in the LSPP. In the case of distributing the load, it may be acceptable for an LSP to traverse through many P node hops, as long as one particular P hop is not much more loaded compared to other P node hops; thus, a percentage or level of tolerance may be established whereby each hop load is within a certain percentage of all other hop loads along the LSP. In the case of minimizing the number of P nodes in the LSPP, the preferred approach is as follows. Given a set of P nodes, and PE nodes PE_{pi} and PE_{pe} , then the edge between P nodes is labeled, with a value of one if bandwidth is available or with a value of infinity if bandwidth is not available. Next, the preferred embodiment runs a known Dijkstra algorithm and queries whether the resultant LSP between PE nodes PE_{pi} and PE_{pe} satisfies delay constraints and has enough bandwidth to carry traffic. These steps are repeated until an optimal LSP is found in both directions between PE nodes PE_{pi} and PE_{pe} . Next, method 20 continues from step 36 to step 38.

[0025] In step 38, central manager CM establishes LSPPs between PE nodes PE_{si} and PE_{se} , that is, for the secondary route. In the preferred embodiment, the selection of these LSPPs is according to the same manner used for the primary route LSPPs described above in connection with step 36 (e.g., the shortest number of P node hops or distributing the load equally or near-equally among all P nodes in the LSP). In addition, however, the P nodes used for the step 36 selection are excluded from consideration in the step 38 selection. This ensures that the backup LSPs are physically disjoint from primary path LSPs while also providing sufficient bandwidth to allow for 1+1 protection. Next, method 20 continues from step 38 to step 40.

[0026] In step 40, central manager CM determines whether the configuration established thus far (e.g., PE node connectivity and primary and backup LSPs) violates any constraint in a set of constraints. In the preferred embodiment, the constraints may include any one or more of the following different constraints, and for those constraints that apply to paths, they are specified in each direction, that is, for each LSP in an LSPP. As a first constraint, a given L2PE node can multi-home only to those PE nodes within a

certain geographical proximity; for instance, this limitation may be on the order of 10 kilometers, by way of example and recognizing that the numbers or range may vary based on implementation. As a second constraint, at a given L2PE node, a VPLS logical entity can connect to only one PE node. As a third constraint, bandwidth into a node also may not exceed any limit provided by the processing capability of that node, that is, different nodes may have different processing capabilities (e.g., PE node PE₁ may process a total of 10G while PE node PE₂ may process only 5G, and so forth; likewise for P nodes). As a fourth constraint, a VPLS comprising of VPLS logical entities within various L2PEs would span or include multiple PE nodes. As a fifth constraint, traffic between a VPLS logical entity and a PE node PE_x should be less than or equal to available bandwidth on the link between that VPLS logical entity and the PE node PE_x. As a sixth constraint, for every VPLS, a PE node can connect to only a subset of certain given number P nodes, that is, for all possible P nodes, some may be excluded as candidate for connectivity to a given VPLS. As a seventh constraint, bandwidth should be equalized in that the sum of bandwidth entering a PE node PE_x (possibly from multiple L2PE nodes) plus the sum of bandwidth entering that PE node PE_x from any P nodes equals the bandwidth leaving that PE node PE_x (towards any one or more L2PE nodes) plus the sum of bandwidth leaving that PE node PE_x towards any P nodes. As an eighth constraint, at a P node, the sum of bandwidth entering the P node must be less than or equal to the sum of the bandwidth capabilities on the output link between that P node and its adjacent P node. As a ninth constraint, the sum of bandwidth entering all PE nodes must equal the sum of bandwidth leaving all PE nodes (to ensure that traffic does not circulate). As a final constraint, end to end delay limits may be provided for one or more paths (e.g., certain paths may carry TDM traffic which is delay sensitive) and these delay limits must be met. Given a set of some or all of these various constraints, step 40 determines whether any constraint in that set has been violated; if so, then method 20 discards the presently-considered set of connections resulting from the immediately-preceding steps 36 and 38 and then returns from step 40 to step 36, so that a new LSPP will be established for the primary configuration followed by the establishment in a successive step 38 of a new LSPP for the backup configuration, which thereafter again are tested by the constraints of step 40.

Eventually, therefore, in many instances step 40 will be satisfied, that is, no constraint will be violated, after which method 20 continues from step 40 to step 42. Note alternatively, however, that at certain instances in the traffic design methodology, there may arise an instance where after numerous loops including steps 36, 38, and 40, there is indeed a point
5 at which no set may be established that satisfies the constraints observed in connection with step 40; toward this end, an appropriate heuristic may be developed by one skilled in the art so as to determine when this point is reached so as to respond appropriately with either a change in flow of method 20, an interrupt, or other desired response so as to continue with the flow while observing that a sufficient number of iterations have
10 occurred so as to likely reach an optimal solution.

[0027] In step 42, central manager CM stores the connectivity information that it has developed in the preceding steps and that satisfied the constraints with respect to step 40; thus, for a case where steps 32 through 40 have been performed a first time with respect to a first VPLS and a first L2PE node, then the associated connectivity is stored. Next,
15 method 20 continues from step 42 to step 44.

[0028] In step 44 central manager CM determines whether the MEN includes another L2PE node that has not yet been considered and that is specified in the information learned in step 30 as desirably including the present VPLS being considered. For example again with reference to Figure 1, if steps 32 through 42 have just been completed a first
20 time with respect to VPLS₁ and relative to L2PE node L2PE₁, then step 44 determines that VPLS₁ also should be considered relative to L2PE node L2PE₃. If an additional L2PE node should thus be considered, then step 44 advances the flow to a step 46 which advances the consideration to that next L2PE node (e.g., L2PE₃), and the flow returns to step 32. If, instead, all L2PE nodes have been considered in relation to the present VPLS (e.g., VPLS₁),
25 then method 20 continues from step 44 to step 48.

[0029] In step 48 central manager CM determines whether the MEN includes another VPLS that has not yet been considered and that is specified in the information learned in step 30. For example again with reference to Figure 1, if steps 32 through 44 have just been

completed a first time with respect to $VPLS_1$, then step 48 determines that $VPLS_2$ also should be considered relative to the MEN. If an additional VPLS should thus be considered, then step 48 advances the flow to a step 50 which advances the consideration to that next VPLS node (e.g., $VPLS_2$), and the flow returns to step 32. If, instead, all VPLSs
5 have been considered, then method 20 continues from step 48 to step 50.

[0030] In step 52, central manager CM determines a cost function of the stored connectivity as relating to all VPLSs across those L2PE nodes to which those VPLSs were deemed to support in the preceding steps 32 through 42. In other words, note that when step 52 is reached, a set of different homing configurations has been developed, where that
10 set includes a different homing configuration for each iteration described above, that is, there may be a different homing configuration for the iteration related to $VPLS_1$ in $L2PE_1$, for the iteration related to $VPLS_1$ in $L2PE_2$, . . . , for the iteration related to $VPLS_1$ in $L2PE_5$, for the iteration related to $VPLS_2$ in $L2PE_1$, for the iteration related to $VPLS_2$ in $L2PE_2$, . . . , for the iteration related to $VPLS_2$ in $L2PE_5$, and so on up to for the iteration related to
15 $VPLS_3$ in $L2PE_5$. Collectively, these iterations and the respective configurations of each therefore provide an entire set of homing configurations, including in this example those in $VPLS_1$ in $L2PE_1$, . . . , $VPLS_3$ in $L2PE_5$. Thus, in step 52, central manager CM determines the cost function for this set of homing configurations. One skilled in the art may determine various different cost functions that may be calculated for step 42. As one
20 example, in one preferred embodiment, the cost function is the total number of VPLS connections (or logical VPLS entities) that are included in the stored connectivity for the MEN system 10. By way of example, if the stored connectivity corresponds to that shown in Figure 1, then step 52 will identify a total number of ten VPLS connections, with three VPLS connections in L2PE node $L2PE_1$, one VPLS connection in L2PE node $L2PE_2$, two
25 VPLS connections in L2PE node $L2PE_3$, three VPLS connections in L2PE node $L2PE_4$, and one VPLS connection in L2PE node $L2PE_5$. Next, method 20 continues from step 52 to step 54.

[0031] In step 54 central manager CM determines whether a sufficient number of iterations have been performed of the above-discussed steps; note that this determination

may be made with respect to the applicable cost function, that is, whether a sufficient diversity in cost functions have been found and that correspond to respective different homing configurations (where recall that a cost function is determined for each homing configuration in the step 52 that operates with respect to that configuration).

5 Alternatively, step 54 may be based on comparing the number of iterations, by having kept a counter of those iterations, with some threshold. Indeed, the overall traffic engineering in a multi-homed VPLS network is presumed by the present inventors to present a so-called NP-Hard problem, that is, one in which there is no absolute best case due to the impossibility of completely considering the preceding alternatives; specifically
10 in the present context, as each connectivity configuration is determined for a VPLS, there is the complexity of evaluating it relative to already-determined connectivity configurations for both that VPLS as well as other VPLSs. Given this attribute, it may be that an absolute number of iterations may be presumed to find an optimal, or at least acceptable, solution. In any event, if the condition(s) of step 54 is such that more iterations
15 are desired, then the flow of method 20 returns continues from step 54 to a step 56 that returns the flow back to the first VPLS (e.g., VPLS1), and the flow returns to step 32; thus, for that instance of step 32 and following, a new set of homing configurations will be ascertained, first for that first VPLS at a first L2PE node, followed by that first VPLS at a second L2PE node, again reaching a consideration of connectivity at all desired L2PE
20 nodes and then for all other VPLSs across all desired L2PE nodes. Alternatively, if the condition(s) of step 54 is such that more iterations are no longer desired, then the flow of method 20 continues from step 54 to step 58.

[0032] In step 58, having been reached after a number of different sets of homing configurations are determined from the preceding steps, then central manager CM
25 chooses the set of homing configurations that has the best cost function among those configurations. For example, in the case where the cost function is number of VPLSs as described above with respect to step 52, then step 58 may select that set of homing configurations which has the largest number of logical VPLS entities for the MEN, and if there is more than one connectivity configuration having a same maximum number of
30 logical VPLS entities, then one skilled in the art may include additional considerations in

the cost function so that one of those connectivity configurations is selected. The best cost function is therefore associated with an output connectivity configuration that includes the number and choice of PE node homes, LSPs, and number of logical VPLS entities. This selection is thus the connectivity that is then to be used in actual configuration of the MEN system 10. Indeed, in the case when central manager CM is part of that system 10, that information is then later used to signal to the various nodes so as to implement this connectivity. Finally, after step 58, method 20 continues to step 60, then method 20 is complete as shown in a complete state 60. Thus, following the complete step 58 or as a part thereof, central manager CM has an optimal connectivity configuration for a multi-homed multiple VPLS system in network 10, and that connectivity configuration may be provided by central manager CM to each node so as to achieve that connectivity.

[0033] The preceding discussion of method 20 and its various steps provide an optimal connectivity configuration for a multi-homed multiple VPLS system as may be implemented when a MEN is being initially configured. However, the preferred embodiments also contemplate that much of method 20 may apply to an existing MEN, where either new traffic and/or new VPLSs are added to that network. Specifically, in such instances, the preferred embodiment determines which PE node to home for a newly-added VPLS and/or possibly adds a new home to a new PE node if blocking is exceeded. In either case, this additional connectivity may be achieved by performing method 20, while eliminating steps 32 and 34; in place of those steps, desired connectivity for the newly-added traffic/VPLS is selected from the already-existing LSPPs while all other steps of method 20 are performed. In this manner, existing traffic and its respective setup is unaffected.

[0034] From the above illustrations and description, one skilled in the art should appreciate that the preferred embodiments provide a method for traffic engineering in a multi-homed VPLS computer network. The preferred embodiments as described provide numerous benefits. As one benefit, the preferred embodiments provide for connectivity with multi-homing, protection and multicast traffic considerations to be included in the optimization process. As another benefit, the preferred embodiments provide dynamic

assignment of homes (i.e., connections) in real time as new VPLSs are established and do not require re-doing the complete optimization problem each time a new VPLS is established. As another benefit, the preferred embodiments do not require re-routing of existing traffic flows as new VPLSs are added. As another benefit, the preferred
5 embodiments are centralized and therefore likely to result in better network resource utilization as compared to a distributed algorithm using routing protocols such as PIM-SM and the like. As another benefit, the preferred embodiments allow QoS considerations to be included in the optimization process. As yet another benefit, the preferred embodiments provide a solution that scales well with an increase in the number of nodes.
10 As a final benefit, while the present embodiments have been described in detail, various substitutions, modifications or alterations could be made to the descriptions set forth above without departing from the inventive scope which is defined by the following claims.